Mike Gibson, Runjie Zhang
ECE 6332 – Calhoun
Design Review 1
8 Oct 2009

Design Review 1

Progress Report

The primary focus of our project-related work up to this point has been planning. Runjie and I have met both with Professor Calhoun and Joe Ryan, a PhD candidate of his, to discuss their work implementing a sub-threshold FPGA system. After a few joint planning sessions, we arrived at a possible project idea: assessing the efficacy of existing leakage reduction strategies in circuits operating in the sub-threshold regime with specific application to sub-threshold FPGA circuits. We also will try to analyze novel novel variations on existing techniques only applicable to circuits operating in the sub-threshold. In particular, we plan to test a minimum of 4 known leakage reduction methods in the sub-threshold regime using Joe Ryan's FPGA circuit.

To this end, Runjie and I have been working to familiarize ourselves with all aspects of this project. Our first goal was to familiarize ourselves with sub-threshold operation in general. To accomplish this, we created a number of simple circuits and studied their simulated behavior, both in super-threshold and sub-threshold regimes. These circuits included an inverter, a nine-stage ring oscillator, and a simple D flip-flop. We were able to identify valid sub-threshold operation, albeit extremely slow operation (on the order of ms clock cycles for the standard ami06 technology in Cadence), for all circuits listed. Additionally, we were able to generate VTC curves both for an inverter in the ami06 technology and an inverter built using 65nm PTM technology. Both curves show valid operation well below the threshold voltage of an individual transistor. Finally, though we ran into some trouble using Ocean due to our lack of experience with the tool, we were able to find and plot the power drawn from the inverter's supply for different sub-threshold on voltages; the plots showed that as sub-threshold voltage decreased, leakage power began to account for significantly more power consumption during normal inverter operation.

In addition to our simulations, we have been reading a number of papers to try and better understand the ideas behind existing leakage models and reduction techniques. We have surveyed six papers, summarized later in this document, to that effect.

Plan for future work

Pre-proposal work

Though Runjie and I have learned much about sub-threshold operation and leakage power already, we have a long way to go before we can successfully contribute to Joe Ryan's project as outlined above. The next major step to understanding leakage reduction will be to create and simulate circuits implementing actual leakage reduction techniques. There are a number of techniques that could potentially be applied to circuits to reduce leakage, including source biasing, taking advantage of the stack effect, multi-threshold CMOS implementations, and, on a larger scale, implementing sleep transistors for inactive logic blocks. To get the most out of this exercise, we will also need to become very familiar with extracting current/energy consumption across a circuit during its operation. To this end, we will also need to become more familiar with the Ocean tools available to us. As we become more familiar with these leakage reduction methods, we will most likely benefit from trying to implement larger, more complex circuits utilizing these techniques. These exercises should provide us with a sufficient understanding of the subject matter to create a reasonable proposal.

Post-proposal work

Once our proposal has been submitted, Runjie and I will need to begin seriously constructing different strategies for implementing leakage reduction techniques in Joe's FPGA circuit. First, we'll need to procure Joe's circuit schematics and become intimately familiar with its operation. This will require some time to study the schematics and run through simulations focusing both on functionality and leakage energy consumption. Once we have sufficiently familiarized ourselves with the FPGA's operation and leakage energy consumption, we'll need to begin implementing our leakage reduction strategies throughout the circuit. As mentioned previously, we're currently planning to implement strategies based on 4 existing leakage reduction methods or novel interpretations of those methods (including things like sub-threshold specific dynamic body biasing). For each strategy, we'll need to perform extensive simulation, both to ensure that our changes have not impacted the functionality of the circuit and to assess the efficacy of our strategy in reducing total leakage energy consumption. Ultimately, we hope to produce a definitive document outlining the applicability of different leakage reduction methods to circuits operating in the sub-threshold regime.

Literature Review

Benton H. Calhoun, Frank A. Honore, and Anantha Chandrakasan.
"A Leakage Reduction Methodology for Distributed MTCMOS".
IEEE Journal of Solid-State Circuits, Vol. 39, No. 5, May, 2004.

This paper presents strategies for minimizing the amount of leakage energy consumed in complex MTCMOS, including FPGAs. The paper addresses a major contributor to leakage energy consumption: sneak leakage paths, or unintended leakage paths that come about as a result of the interactions between sections of a circuit. It also proposes design strategies for local sleep devices (as opposed to global sleep devices) intended to reduce both standby leakage and sneak leakage paths, ultimately outlining 4 major rules for preventing sneak leakage. These rules include: MTCMOS gates that share output with high-V_T elements should use both polarity sleep devices; MTCMOS gates sharing outputs with other MTCMOS gates must have the same polarity sleep devices; MTCMOS gates sharing outputs with gates having both polarity sleep devices must themselves have both polarity sleep devices; and, finally, sleep devices should not be shared between lines that connect outputs from multiple high-V_T elements.

The paper culminates in the simulation of an FPGA architecture utilizing these strategies. It briefly outlines the applicability of local sleep devices to FPGA circuits. Overall, localized sleep regions and actively designing to reduce sneak leakage paths reduces leakage energy consumption by ~8x when the entire circuit is in sleep mode, and by about 2.2x for CLBs in different configurations.

Kao, J.; Narendra, S.; Chandrakasan, A., "Subthreshold leakage modeling and reduction techniques [IC CAD tools]," *Computer Aided Design, 2002. ICCAD 2002. IEEE/ACM International Conference on* , vol., no., pp. 141-148, 10-14 Nov. 2002

This paper will most likely be extremely useful in our work for this project. The paper begins with a simple model for sub-threshold leakage current taking into consideration a number of factors including channel length and other V_T variation effects. This provides an improvement over previous leakage current models which only tried to determine upper and lower bounds for leakage current (according to the paper, leakage current in micrometer technologies typically had negligible leakage current contributions to overall power consumption, therefore accurate leakage current estimations

were not necessarily needed when determining the energy consumption of a circuit).

After providing additional recommendations for CAD tool requirements pertaining to leakage simulation, the paper outlines a number of leakage reduction strategies, including: source biasing, utilizing the stack effect, dual V_T partitioning, MTCMOS implementations (including the utilization of sleep transistors in larger networks), and variable threshold CMOS devices. The paper concludes with a brief analysis of optimal V_{DD}/V_T operating points, and a reiteration of the importance of accurate sub-threshold leakage current models for implementing future nm-scale technologies.

Kumar, A.; Anis, M., "Dual-Vt design of FPGAs for subthreshold leakage tolerance," *Quality Electronic Design, 2006. ISQED '06. 7th International Symposium on* , vol., no., pp.6 pp.-740, 27-29 March 2006

This paper proposes a dual-V_T implementation of a SRAM-based FPGA architecture (consisting of SRAM LUTs that comprise basic logic elements and an island-based routing architecture) in an attempt to mitigate perceived shortcomings of strategies such as those outlined in the previous summary and other works (including difficulty of implementation of source biasing, limited application of local sleep device logic, etc). To that end, the paper presents a proposed FPGA architecture, a CAD flow for the proposed FPGA, and algorithms associated with the CAD flow to aid in the creation of dual-V_T FPGA circuits.

Specifically, the paper outlines a strategy for the assignment of high-V_T devices to a large percentage of the logic blocks within the CLBs (their justification for this includes the fact that routing delays are much greater than logic delays). This strategy is implemented in the paper's proposed 6-stage CAD flow (developed using VPR and T-V pack tools). In summary, the CAD flow begins by assuming all devices are low-V_T devices. The algorithm progresses as follows: first, a delay estimation of the various paths on the circuit is made; then high-V_T subblocks are assigned based on these estimations; the subblocks are re-clustered throughout the circuit; the new blocks are placed and routed; and finally, the new leakage power savings are computed. The paper continues with an analysis of the possible architectures that result from this algorithm, including heterogeneous architectures (type1 and type2), homogeneous architectures (type1), and homogeneous architectures (type2). Finally, the paper concludes with a summary of its evaluation methods and an example implementation; in the example implementation, average power savings of 40% were observed for a circuit with a 10% active time and a cluster size of 12 blocks.

Mingoo Seok; Hanson, S.; Yu-Shiang Lin; Zhiyoong Foo; Daeyeon Kim; Yoonmyung Lee; Nurrachman Liu; Sylvester, D.; Blaauw, D., "The Phoenix Processor: A 30pW platform for sensor applications," *VLSI Circuits, 2008 IEEE Symposium on* , vol., no., pp.188-189, 18-20 June 2008

In this short paper, an implementation of a low power processor for sensor applications is presented. The processor occupies an area of only $1mm^3$ when coupled with an on-die battery, and can operate for over 1 year in its default configuration. It switches between sleep mode and active mode to save power. Unlike the traditional power gating methods which utilize HVT(high-V_T), the Phoenix Processor uses MVT(medium-V_T), which helps the processor focus more on power efficiency than performance. The paper investigates the relationship between footer width and frequency/energy in both sleep and active mode, and finds that total energy decreases by 4 orders of magnitude with a footer and decreases by 2.5x with a narrow order (0.66um compared to 28um). The remainder of the paper discusses the architecture and algorithm level design of the processor in detail.

Benton H. Calhoun*, Alice Wang, and Anantha P. Chandrakasan.
"Modeling and Sizing for Minimum Energy Operation in Subthreshold Circuits".
IEEE Journal of Solid-State Circuits, Vol. 40, No. 9, pages 1778-1786, September, 2005.

This paper presents a model for energy savings in sub-threshold circuits. It provides an equation for total energy consumption and shows that if V_{DD} decreases, dynamic power decreases as well, thus decreasing total energy consumed in the circuit up to a point. Deep in sub-threshold operation, leakage energy begins to increase exponentially, eclipsing dynamic energy consumption and leading to an increase in total energy. The paper proposes an equation to compute the minimum energy point $V_{DD\text{-}opt}$ and shows that this point is independent of V_T.

Transistor sizing also has an effect on operation in sub-threshold regimes. Minimum V_{DD} operation occurs when the PMOS and NMOS devices have the same current. Also, sizing not only affects the energy consumption in a circuit, but can also lead to faulty operation if not undertaken with care.

In the remainder of this paper, the author also shows that the minimum energy point also depends on the operating conditions such as temperature (to a lesser extent), duty cycle and total workload.

Joseph F. Ryan, Jiajing Wang, and Benton H. Calhoun
"Analyzing and Modeling Process Balance for Sub-threshold Circuit Design" (pdf). | Slides (pdf)
GLSVLSI, pp. 275-280, March 2007.

This paper begins by asserting that process variation is the most daunting obstacle to efficient sub-threshold circuit design. In the sub-threshold region, device symmetry is quite important because it provides the best-case for circuit stability and affects the robustness of static CMOS circuits. The author continues with a definition for sub-threshold process balance: it is the relative strength of PMOS and NMOS devices in a process. A process is balanced if the NMOS and PMOS currents are the same during switching. The process balance is affected by several factors, including differences in threshold voltage, differences in sub-threshold slope, drain-induced barrier lowering (DIBL), process scaling and transistor sizing.

Static CMOS is the most common logic style used in sub-threshold circuits due to its robustness. Process balance will alter the robustness of static logic by varying noise margins. Additionally, process imbalance can have an impact on the static noise margin (SNM) of a SRAM circuit. To obtain a higher SNM, the imbalanced processes should be adjusted (through resizing, for example) towards the more balanced process.

In the rest of the paper, the author develops a linear model to predict the VM of a minimum-sized inverter to within a mean 3.2% error for a PTM process.

Simulation Results

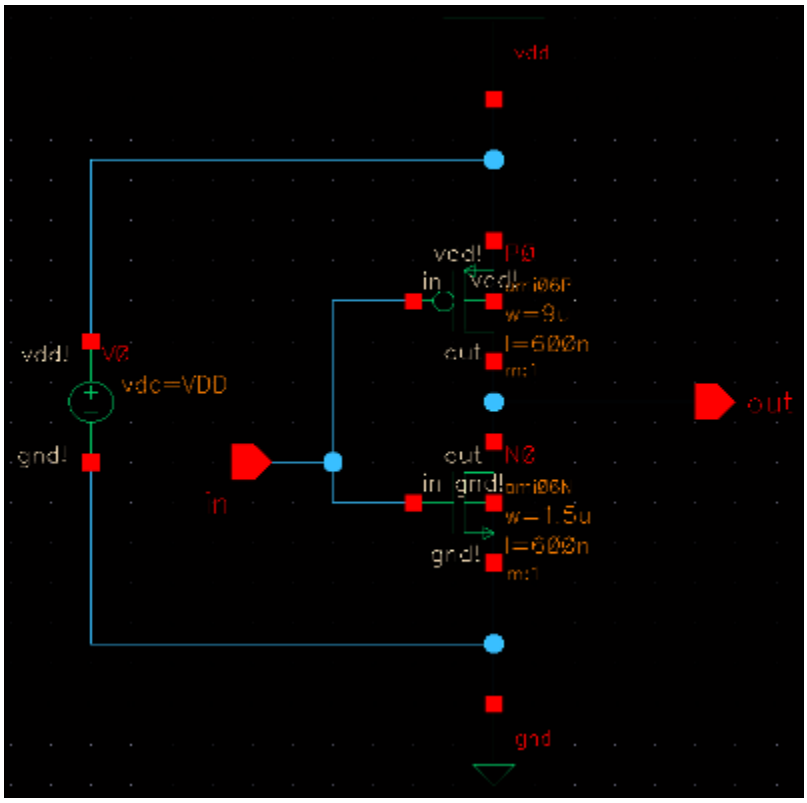

Figure 1. Implementation of a simple inverter in ami06 technology.

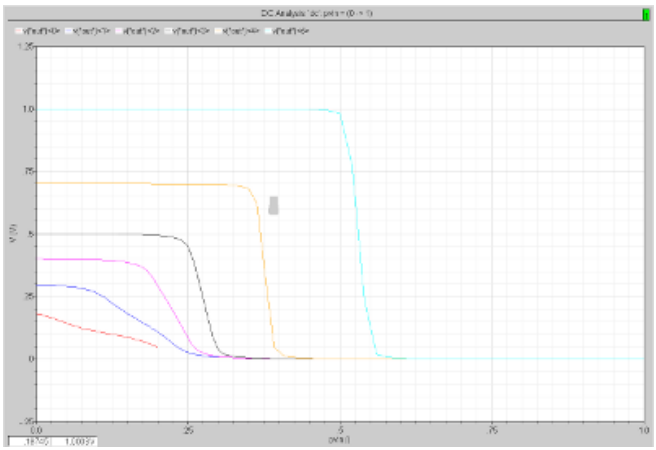

Figure 2. VTC of inverter implemented above at various supply voltages (increasing from bottom to top, 0.2 to 1.0V, Vout vs. Vin). Note that noise margins remain acceptable in deep sub-threshold operation.

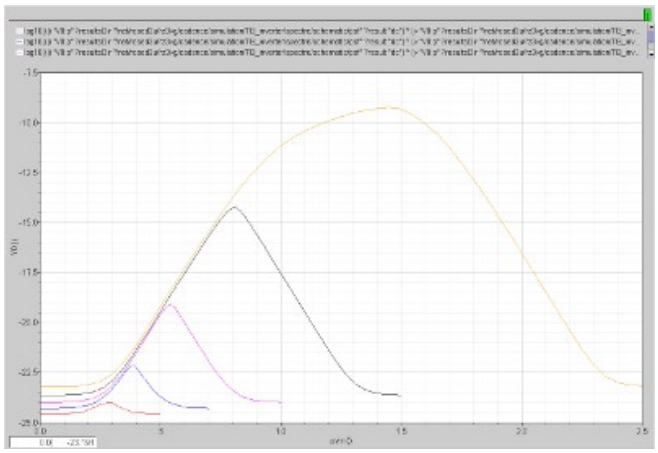

Figure 3. Total power drawn from the supply during a DC sweep of input at various supply voltages (increasing from bottom to top, 0.2 – 1.0 V, log(Id) vs. Vin). Note that leakage power in cutoff state tends toward the same order of magnitude as dynamic power at lower supply voltages.

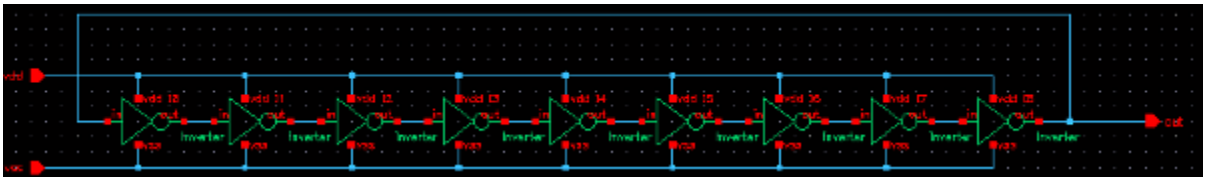

Figure 4. A simple nine-stage ring oscillator based on the inverter structure demonstrated above.

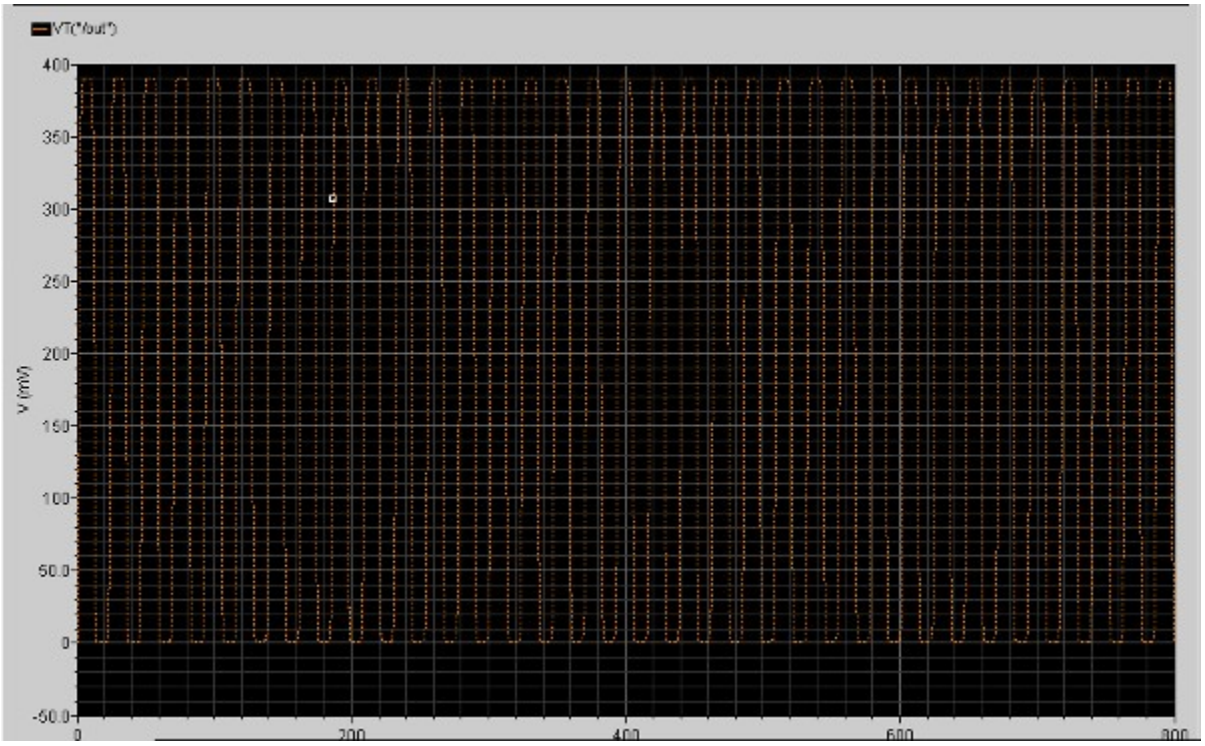

Figure 5. The sub-threshold transient response (Vdd = 400mV) of the ring oscillator. Note that the time axis is in ms.

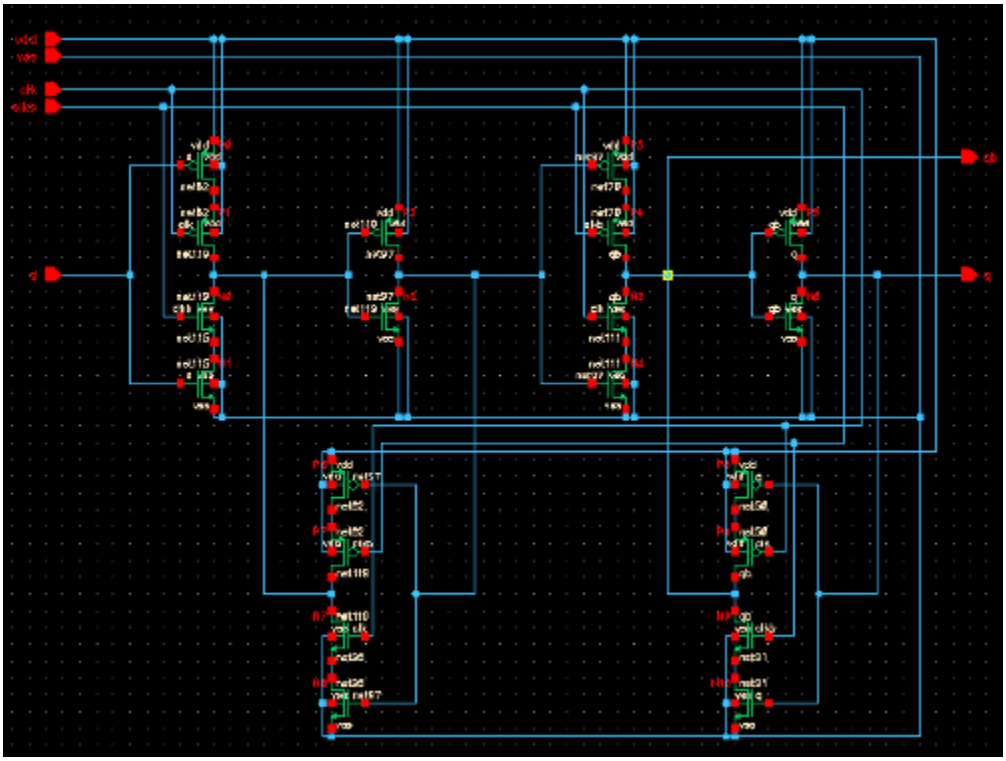

Figure 6. A D flip-flop implemented in 65nm PTM technology.

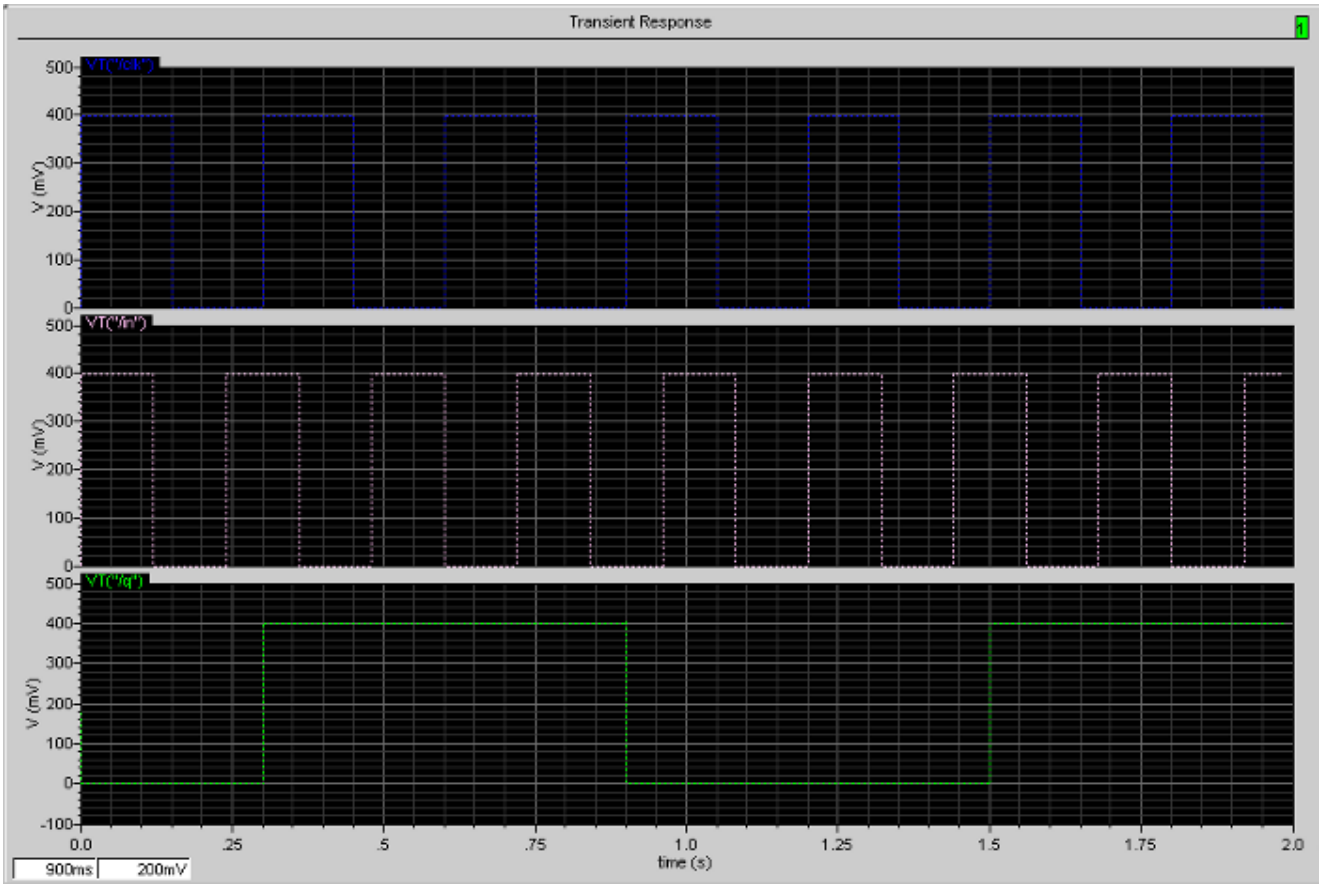

Figure 7. The sub-threshold transient response of the D flip-flop. Note Vdd is 400mV.